

## Amendments to the Claims

1. (Original): A method for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising the steps of:

- monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and
- predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.

2. (Original): The method of Claim 1, wherein said predicting step further comprises the steps of:

- computing a regression model of prediction parameters on the number of active devices;
- constructing an empirical distribution of the number of active devices; and
- combining the regression model and the empirical distribution to produce a prediction model.

3. (Original): The method of Claim 2, wherein said step of combining the regression model and the empirical distribution comprises the step of computing, with respect to the empirical distribution, an expected value of each of one or more of the prediction parameters.

4. (Original): The method of Claim 3, wherein said combining step further comprises the steps of:

- for each of the one or more prediction parameters,
- for each of the monitored values of the number of active devices, computing confidence intervals for the one or more prediction parameters; and
- selecting a corresponding one of the confidence intervals for the expected value of each of the one or more prediction parameters.

5. (Original): The method of Claim 3, wherein said combining step further comprises the steps of:

computing confidence intervals for the prediction parameters, for each of the monitored values of the number of active devices; and

combining the confidence intervals with the empirical distribution, to produce modified confidence intervals for the prediction parameters.

6. (Original): The method of Claim 1, further comprising the step of managing a resource capacity of the computer system, based upon the predicted subsequent resource utilization

7. (Original): The method of Claim 1, wherein the computer system comprises computer software, and said method further comprising the step of rejuvenating the computer software, based upon the predicted subsequent resource utilization.

8. (Original): The method of Claim 1, further comprising the step of dynamically allocating at least one resource of the computer system, based upon the predicted subsequent resource utilization:

9. (Original): The method of Claim 1, wherein said predicting step comprises the steps of:

identifying any of the plurality of devices that are relevant to a monitored resource; and

restricting at least one subsequent operation of the computer system that corresponds to the monitored resource to use only devices identified as relevant to the monitored resource from among the plurality of devices.

10. (Original): The method of Claim 9, wherein said identifying step comprises the steps of:

for a given device currently being evaluated for relevance to the monitored resource,

computing a first regression model for the monitored resource on the number of active devices, based upon data acquired when the given device is active;  
computing a second regression model for the monitored resource on the number of active devices, based upon data acquired when the given device is inactive; and  
comparing the first and the second regression models to determine whether the given device is relevant to the monitored resource.

11. (Original): The method of Claim 10, wherein said comparing step comprises the step of determining whether the first and the second regression models are statistically equivalent for a same number of active devices other than the given device.

12. (Original): The method of Claim 1, wherein said predicting step comprises the steps of:

dividing the plurality of devices into device classes; and  
counting the number of active devices in each of the device classes.

13. (Original): The method of Claim 12, further comprising the step of fitting a prediction model for a monitored resource, wherein prediction model parameters of the prediction model depend on the number of active devices in each of the device classes.

14. (Original): The method of Claim 12, further comprising the step of computing a prediction model for the number of active devices in each of the device classes.

15. (Original): The method of Claim 1, wherein said predicting step comprises the steps of:

computing a regression model of prediction parameters on the number of active devices;  
constructing a prediction model for a distribution of the number of active devices;  
and  
combining the regression model and the prediction model for the distribution of active devices to produce a prediction model for one or more of the prediction parameters.

16. (Original): The method of Claim 15, wherein said predicting step further comprises the steps of:

- computing a regression model of prediction parameters on the number of active devices;

- constructing an empirical distribution of the number of active devices; and

- combining the regression model and the empirical distribution to produce a prediction model.

17. (Original): The method of Claim 16, wherein said step of combining the regression model and the empirical distribution comprises the step of computing an expected value of the prediction parameters with respect to the empirical distribution.

18. (Original): The method of Claim 17, wherein said combining step further comprises the steps of:

- for each of the one or more prediction parameters,

- for each of the monitored values of the number of active devices, computing confidence intervals for the one or more prediction parameters; and

- selecting a corresponding one of the confidence intervals for the expected value of each of the one or more prediction parameters.

19. (Original): The method of Claim 17, wherein said combining step further comprises the steps of:

- computing confidence intervals for the prediction parameters for each of the monitored values of the number of active devices; and

- combining the confidence intervals for the prediction parameters for each of the monitored values of the number of active devices with the empirical distribution of the number of active devices, to produce modified confidence intervals for the prediction parameters.

20. (Original): The method according to Claim 1, wherein said method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

21. (Original): A method for identifying resource saturation in a computer system having a plurality of devices, comprising the steps of:

monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.

22. (Original): The method of Claim 21, wherein said identifying step comprises the steps of:

fitting a regression model of a monitored resource on the number of active devices; and

detecting, in the regression model, departures of the monitored resource from linearity.

23. (Original): The method of Claim 21, wherein said identifying step comprises the steps of:

fitting a regression model of prediction parameters of the monitored resource on the number of active devices; and

detecting departures from linearity of the prediction parameters of the monitored resource.

24. (Original): The method of Claim 21, wherein said identifying step comprises the steps of:

identifying any of the plurality of devices that are relevant to a monitored resource; and

restricting at least one subsequent operation of the computer system that corresponds to the monitored resource to use only devices identified as relevant to the monitored resource from among the plurality of devices.

25. (Original): The method of Claim 24, wherein said identifying step comprises the steps of:

for a given device currently being evaluated for relevance to the monitored resource,

computing a first regression model for the monitored resource on the number of active devices, based upon data acquired when the given device is active;

computing a second regression model for the monitored resource on the number of active devices, based upon data acquired when the given device is inactive; and

comparing the first and the second regression models to determine whether the given device is relevant to the monitored resource.

26. (Original): The method of Claim 24, wherein said comparing step comprises the step of determining whether the first and the second regression models are statistically equivalent for a same number of active devices other than the given device.

27. (Original): The method of Claim 21, wherein said identifying step comprises the steps of:

dividing the plurality of devices into device classes; and

counting the number of active devices in each of the device classes.

28. (Original): The method of Claim 27, further comprising the step of fitting a prediction model for a monitored resource, wherein prediction model parameters of the prediction model depend on the number of active devices in each of the device classes.

29. (Original): The method of Claim 27, further comprising the step of computing a prediction model for the number of active devices in each of the device classes.

30. (Original): The method according to Claim 21, wherein said method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

31. (Original): A method for predicting effects of adding a new device on a computer system having a plurality of devices, the method comprising the steps of:  
monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and  
predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.

32. (Original): The method of Claim 31, wherein said predicting step comprises the steps of:  
selecting a monitored resource; and  
predicting the effects of adding the new device with respect to the selected monitored resource, based upon the monitored values of the resource utilization and the number of active devices.

33. (Original): The method of Claim 32, wherein said predicting step comprises the steps of:  
constructing a first prediction model of a distribution of the number of active devices; and  
modifying the first prediction model to produce a modified prediction model of the distribution of the number of active devices that accounts for the new device.

34. (Original): The method of Claim 33, further comprising the steps of:  
computing a first prediction model for the selected monitored resource, based upon the first prediction model of the distribution of the number of active devices;  
producing a first prediction of the selected monitored resource using the first prediction model for the selected monitored resource;

computing a modified prediction model for the selected monitored resource to account for the new device, based upon the modified prediction model of the distribution of the number of active devices;

producing a second prediction of the selected monitored resource, based upon the modified prediction model for the selected monitored resource; and

comparing the first prediction and the second prediction of the selected monitored resource to evaluate the effects of adding the new device.

35. (Original): The method of Claim 31, wherein said predicting step comprises the step of:

dividing the plurality of devices into device classes; and

counting the number of active devices in each of the device classes.

36. (Original): The method of Claim 32, further comprising the steps of:

identifying any of the plurality of devices that are relevant to a monitored resource; and

restricting at least one subsequent operation of the computer system that corresponds to the monitored resource to use only devices identified as relevant to the monitored resource from among the plurality of devices.

37. (Original): The method according to Claim 31, wherein said method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

38. (Original): An apparatus for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising:

a monitoring device for monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and



a prediction device for predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.

39. (Original): An apparatus for identifying resource saturation in a computer system having a plurality of devices, comprising:

a monitoring device for monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

a forecasting device for identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.

40. (Original): An apparatus for predicting effects of adding a new device on a computer system having a plurality of devices, the apparatus comprising:

a monitoring device for monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

a forecasting device for predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.

41. (New): A method for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising the steps of:

monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and

predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices;

wherein said predicting step further comprises the steps of

computing a regression model of prediction parameters on the number of active devices,

constructing an empirical distribution of the number of active devices, and

combining the regression model and the empirical distribution to produce a prediction model.